

A SPATIAL ANALYSIS OF MULTIVARIATE OUTPUT FROM REGIONAL CLIMATE MODELS

BY STEPHAN R. SAIN¹, REINHARD FURRER² AND NOEL CRESSIE³

*National Center for Atmospheric Research, University of Zurich and
Ohio State University*

Climate models have become an important tool in the study of climate and climate change, and ensemble experiments consisting of multiple climate-model runs are used in studying and quantifying the uncertainty in climate-model output. However, there are often only a limited number of model runs available for a particular experiment, and one of the statistical challenges is to characterize the distribution of the model output. To that end, we have developed a multivariate hierarchical approach, at the heart of which is a new representation of a multivariate Markov random field. This approach allows for flexible modeling of the multivariate spatial dependencies, including the cross-dependencies between variables. We demonstrate this statistical model on an ensemble arising from a regional-climate-model experiment over the western United States, and we focus on the projected change in seasonal temperature and precipitation over the next 50 years.

1. Introduction. Many processes in the Earth system cannot be directly observed, and computer modeling has become a primary mode for studying these processes. These computer models often encapsulate entire fields of knowledge, providing a virtual laboratory for understanding physical relationships and serving as a basis for making predictions. The Earth's climate, for example, is determined by the flows of energy, water, gases, etc., within and between the different components of the climate system, including atmosphere, oceans, terrestrial and marine biospheres, sea ice, etc. Climate

Received October 2009; revised June 2010.

¹Supported in part by NSF Grants DMS-07-07069 and ATM-05-34173.

²Supported in part by NSF Grant DMS-06-21118. In addition, much of the research in this paper was done while the second author was an assistant professor in the Department of Mathematical and Computer Sciences at the Colorado School of Mines in Golden, CO.

³Supported in part by the Office of Naval Research under Award N00014-08-1-0464.

Key words and phrases. Lattice data, Markov random field (MRF), conditional autoregressive (CAR) model, Bayesian hierarchical model, climate change.

This is an electronic reprint of the original article published by the [Institute of Mathematical Statistics](#) in *The Annals of Applied Statistics*, 2011, Vol. 5, No. 1, 150–175. This reprint differs from the original in pagination and typographic detail.

models attempt to represent this system, as well as to incorporate anthropogenic forcings to assess the impact of human intervention.

While climate models are deterministic, the output generated by these models is complex and subject to a number of sources of uncertainty. Initial climate states, assumptions about future forcings, and, of course, our understanding (or lack thereof) of the physical processes and their representation in computer models, are all issues that may lead to uncertainty in the model output. To gain a better understanding of this model uncertainty, there is an increasing use of ensembles consisting of multiple model runs. These experiments may involve varying initial conditions (simple ensembles), model physics (perturbed-physics ensembles), specific models (multi-model ensembles), or some combination of all three, in an attempt to capture the range of variation in the model output.

Climate is often considered the long-term average (or, more generally, the long-term distribution) of weather, and climate models typically simulate decades of weather to capture this long-term behavior. These simulations can take weeks or even months of computing time on high-performance supercomputers. Furthermore, as interest continues to grow in climate models with higher spatial resolution, in particular, to study regional and local impacts of climate and climate change, the computational demands of climate modeling continue to grow. Even with the increased use of ensembles, the number of ensemble members is typically limited due to constraints on the models available, computation, funding, etc. Hence, statistical methods are necessary to quantify the distribution and breadth of variation of the model output in the ensemble. To this end, we introduce a hierarchical statistical model to capture the multivariate spatial distribution of the output fields (e.g., the joint spatial distribution of temperature and precipitation) from a regional-climate-model ensemble. With this statistical representation of the model output, we can present probabilistic projections of regional climate change based on the ensemble. Furthermore, by considering multiple output fields simultaneously, we can incorporate correlations between these fields, improving joint projections necessary for many climate-impacts studies (e.g., agriculture, water management, public health, etc.).

1.1. Regional climate models. The climate of a region is determined by processes that exist at planetary, regional, and local spatial scales and across a wide range of temporal scales (multi-decadal to sub-daily). This creates serious difficulties when attempting to construct computer models that can simulate regional climate. Atmosphere-ocean general circulation models (AOGCMs or, more simply, GCMs) couple an atmospheric model with an ocean model and seek to simulate the Earth's global climate system. Due to model complexity, the need to simulate climate over decadal

and even centennial time scales, and computational limitations, these models typically have grid boxes on the spatial scale of 200 to 500 km. While these models are extremely useful for investigating the large-scale circulation and forcings that affect the Earth’s global climate, there are limitations to their use for regional and local projections that might be of interest to the climate-impacts community.

Recognizing the need to include large-scale processes, even when studying regional climate, as well as the ability of GCMs to capture such phenomena, there is considerable attention on developing downscaling methodologies. Downscaling refers simply to generating information on the basis of a GCM, but at spatial scales below that of the GCM. There are two main types of downscaling, dynamical and statistical. Statistical downscaling is a computationally efficient approach that uses empirical relationships to connect the coarse-resolution GCM output to regional and local variables. This approach needs fine-scale and long-time-duration observational data, and there is some uncertainty about the stability of these empirical relationships over long periods of time, especially with varying forcings (e.g., increasing CO₂ concentrations).

Dynamical downscaling involves using high-resolution climate models. Of course, there are limitations due to computational demands and a price to be paid for the higher resolution. One approach uses only the atmospheric component of a GCM with, for example, observed ocean temperatures. These so-called time-slice experiments are globally consistent, but they generally use many of the same formulations as the coarse-scale GCMs.

Regional climate models (RCMs) are the focus of this research and are another dynamical approach based on high-resolution climate models. These models typically focus on a limited spatial domain, have grid boxes on the scale of 20 to 100 km, and there are often simplifications of ocean processes in these models. They also use initial conditions and time-dependent lateral boundary conditions from the GCM (e.g., winds, temperature, and moisture). Hence, global circulation and large-scale forcings are consistent with the GCM, but, with the higher-resolution forcings included (e.g., topography, land cover, etc.), these models have the potential to actually enhance the simulation of climate on regional and local scales. Of course, RCMs can be influenced by potential biases in the GCM, and there is a lack of two-way interactions between the driving GCM and the RCM.

For those interested in understanding more about climate and climate change and global and regional climate modeling, we refer them to the Intergovernmental Panel on Climate Change (IPCC, <http://www.ipcc.ch>) assessment reports, and, in particular, to the contributions of Working Group I to the Third Assessment Report [Houghton et al. (2001)] and the Fourth Assessment Report [Solomon et al. (2007)]. These documents not only include excellent overviews of the issues but also numerous scientific references for more in-depth coverage.

1.2. *A statistical representation of climate-model output.* We propose a statistical model for combining the output from simple ensembles of RCMs in order to characterize the distribution of the model output. This statistical model will be formulated through what has now become the standard three-level hierarchical formulation, namely, data model, process model, and parameter model (prior distribution). The data model links the RCM output to an unobserved spatial process, where this process model is formulated to capture the spatial variation in the RCM output. Both the data model and the process model depend on unknown parameters to which a prior distribution is assigned.

This basic hierarchical approach has been used in other settings for combining climate-model output. For example, Tebaldi et al. (2005) focuses on univariate summaries of temperature change from an ensemble of GCMs. Smith et al. (2009) also explore univariate summaries of temperature change from an ensemble of GCMs and link these summaries over different regions around the globe. Tebaldi and Sansó (2009) extend these approaches to bivariate models of temperatures and precipitation. Furrer et al. (2007a, 2007b) study univariate spatial summaries and Berliner and Kim (2008) study bivariate time series, again constructed from an ensemble of GCMs. More recently, Cooley and Sain (2010) and Kang, Cressie and Sain (2010) have applied hierarchical models to ensembles of RCMs. However, to our knowledge, this paper is the first approach of this kind for a *spatial* analysis of *multivariate* output from RCMs.

At the heart of this statistical model is an implementation of a multivariate Markov random field (MRF) for lattice data that offers great flexibility in modeling the spatial cross-dependencies between variables. We emphasize this capability in this setting, as the underlying physical behavior of the climate system suggests the potential for significant spatial cross-dependencies, in particular, for key variables like *temperature* and *precipitation*. While these two output fields are the focus of this work, we note that the basic approach based on the multivariate MRF presented here can be easily modified to consider other output fields as well.

MRF models are also excellent tools for analyzing data laid out on regular spatial lattices, such as those associated with images, remote-sensing, climate models, etc., or on irregular spatial lattices, such as U.S. census divisions (counties, tracts, or block-groups) or other administrative units. In contrast to geostatistical methods that model spatial dependence through the specification of a covariance function (typically based on distances between spatial locations), Gaussian MRF models represent the conditional expectation of an observation at a spatial location as a linear combination of observations at neighboring locations. Spatial dependence is induced through this conditional autoregression and the choice of neighborhoods.

In addition, using a MRF formulation will allow us to incorporate computational advantages due to the gridded nature of the climate-model output and the sparseness that is characteristic of the spatial-precision matrices (inverse covariance matrices) that are specified in such models. Furthermore, the multivariate nature of the statistical model (more than one model output considered at each grid box) will allow for more complex inferences that are of use to those studying impacts of climate and climate change.

1.3. Outline. In Section 2 an overview of MRF models is presented, followed by a description of the new formulation in Section 3. Section 4 contains the details of our hierarchical specification. An extensive study of an application using a simple ensemble of RCM output, focusing on changes in seasonal temperature and precipitation, will be presented in Section 5. Concluding remarks are given in Section 6.

2. MRF and CAR models. Besag (1974) laid out the basic framework for MRF models. For random variables y_1, \dots, y_n observed at n locations on a spatial-lattice structure, the collection of conditional distributions $f(y_i|y_{-i})$, $i = 1, \dots, n$ (where y_{-i} refers to all random variables except the i th one), can be combined under certain regularity conditions to form a joint distribution $f(y_1, \dots, y_n)$. Rue and Held (2005) can be consulted for an excellent exposition of the theory of MRFs; see also the reviews in the texts by Cressie (1993), Banerjee, Carlin and Gelfand (2004), and Schabenberger and Gotway (2005). Conditional autoregressive (CAR) models are special cases of MRF models where the conditional distributions are assumed to be Gaussian.

2.1. Univariate CAR models. In the univariate setting and assuming Gaussian conditional distributions for $f(y_i|y_{-i})$, the conditional mean and conditional variance associated with $f(y_i|y_{-i})$ are specified as

$$E[y_i|y_{-i}] = \mu_i + \sum_{j \neq i}^n b_{ij}(y_j - \mu_j) \quad \text{and} \quad \text{Var}[y_i|y_{-i}] = \tau_i^2 > 0,$$

where $b_{ii} = 0$; $i = 1, \dots, n$. Under regularity conditions, this collection of conditional distributions gives rise to a joint Gaussian distribution,

$$(1) \quad \mathcal{N}(\boldsymbol{\mu}, (\mathbf{I} - \mathbf{B})^{-1}\mathbf{T}),$$

where $\boldsymbol{\mu} = (\mu_1, \dots, \mu_n)'$, \mathbf{I} is an $n \times n$ identity matrix, \mathbf{B} is the $n \times n$ matrix with the (i, j) th element b_{ij} , and $\mathbf{T} = \text{diag}(\tau_1^2, \dots, \tau_n^2)$. The regularity conditions are on the spatial-dependence parameters, $\{b_{ij}\}$, and they ensure that the resulting matrix, $(\mathbf{I} - \mathbf{B})^{-1}\mathbf{T}$, is a *bona fide* covariance matrix; that is, $(\mathbf{I} - \mathbf{B})^{-1}\mathbf{T}$ is symmetric and positive-definite. The spatial dependence is induced by the autoregression, which is determined by setting the coefficients $b_{ij} \neq 0$ if $j \in N_i$ (and 0 otherwise), where N_i is a collection of indices that define a neighborhood of the i th location in the spatial lattice.

2.2. Multivariate CAR models. Mardia (1988) extended the MRF model of Besag (1974) to the multivariate setting where there is more than one measurement at each lattice point. In particular, let \mathbf{y}_i be a p -dimensional random vector. Then, for $i = 1, \dots, n$, let $f(\mathbf{y}_i|\mathbf{y}_{-i})$ be a Gaussian distribution of \mathbf{y}_i , given all random vectors except the i th, with

$$E[\mathbf{y}_i|\mathbf{y}_{-i}] = \boldsymbol{\mu}_i + \sum_{j \neq i} \mathbf{B}_{ij}(\mathbf{y}_j - \boldsymbol{\mu}_j) \quad \text{and} \quad \text{Var}[\mathbf{y}_i|\mathbf{y}_{-i}] = \mathbf{T}_i,$$

where $\boldsymbol{\mu}_i$ is a p -dimensional vector, \mathbf{B}_{ij} is a $p \times p$ matrix, and \mathbf{T}_i is a $p \times p$ covariance matrix. Assume that $\mathbf{B}_{ij}\mathbf{T}_j = \mathbf{T}_i\mathbf{B}'_{ji}$, for all $i, j = 1, \dots, n$ (to ensure symmetry). Further, assume $\mathbf{B}_{ii} = -\mathbf{I}$ and $\mathbf{B}_{ij} \neq \mathbf{0}$ for $j \in N_i$ and $i = 1, \dots, n$. Under the assumption that the $np \times np$ matrix $\text{Block}(-\mathbf{B}_{ij})$ (i.e., a block matrix with blocks given by $-\mathbf{B}_{ij}$) is positive-definite, Mardia (1988) establishes that $\mathbf{y} = (\mathbf{y}'_1, \dots, \mathbf{y}'_n)'$ follows a $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ distribution where $\boldsymbol{\mu} = (\boldsymbol{\mu}'_1, \dots, \boldsymbol{\mu}'_n)'$ and $\boldsymbol{\Sigma} = (\text{Block}(-\mathbf{T}_i^{-1}\mathbf{B}_{ij}))^{-1}$.

As written, this formulation is overparameterized and there have been a number of efforts in the literature that focus on ways of specifying the parameters in the basic model. See, for example, Billheimer et al. (1997), Kim, Sun and Tsutakawa (2001), Pettitt, Weir and Hart (2002), Carlin and Banerjee (2003), Gelfand and Vounatsou (2003), Jin, Carlin and Banerjee (2005), Jin, Banerjee and Carlin (2007), Daniels, Zhou and Zou (2006), Sain and Cressie (2007), among others.

3. An alternative formulation of a multivariate MRF. It is our experience that the basic multivariate MRF model of Mardia (1988) is difficult to implement in practice without dramatic simplification of the matrices representing the spatial dependence parameters [e.g., Billheimer et al. (1997)] or the use of restrictive priors on the elements of these same matrices [Sain and Cressie (2007)]. Here, we fully develop a new way of representing multivariate lattice data that was suggested, but not implemented, by Sain and Cressie (2007), for the purposes of analyzing an RCM experiment.

The multivariate extension of the framework laid out by Besag (1974) and explored by Mardia (1988) is based on the assumption of a multivariate observation at each point on a standard two-dimensional spatial lattice. Fundamental to our approach is thinking of multivariate lattice data as univariate data on a more complex lattice structure.

In particular, this more complex lattice structure is conceptualized as a “stacking” of the lattices associated with each variable. Neighborhoods are defined by connections between locations for each variable within a lattice and again for locations across each lattice structure. A two-dimensional example of these different types of neighborhoods is shown in Figure 1. A within-variable dependence structure is induced by connecting locations

within a lattice associated with a particular variable (left frame). Cross-dependencies, both within a location (middle frame) and across locations (right frame), are induced through connections between the lattices for different variables.

The key feature of this approach is that it still falls within the original univariate framework of Besag (1974) outlined in Section 2.1. Let y_{ij} denote the j th variable observed at the i th location on the lattice. Then, for each Gaussian conditional distribution, the mean and variance need to be specified. With sums on the right-hand side corresponding to specific types of neighborhoods in Figure 1, the conditional mean is given by

$$(2) \quad \begin{aligned} E[y_{ij}|y_{-\{ij\}}] = & \mu_{ij} + \sum_{k \neq i} b_{ijkj}(y_{kj} - \mu_{kj}) \\ & + \sum_{\ell \neq j} b_{ij\ell}(y_{i\ell} - \mu_{i\ell}) + \sum_{k, \ell \neq i, j} b_{ijk\ell}(y_{k\ell} - \mu_{k\ell}) \end{aligned}$$

with the conditional variance given by $\text{Var}[y_{ij}|y_{-\{ij\}}] = \tau_{ij}^2$, for all lattice points $i = 1, \dots, n$, variables $j = 1, \dots, p$, and with $y_{-\{ij\}}$ denoting all components of \mathbf{y} except for the (i, j) th.

In the conditional mean, the coefficients in the first summation, $\{b_{ijkj}; i = 1, \dots, n, k \in N_i, j = 1, \dots, p\}$, represent connections within a particular layer and control conditional dependence between the i th lattice point and neighboring points for the j th variable (left frame in Figure 1). The coefficients in the second summation, $\{b_{ij\ell}; i = 1, \dots, n, j \neq \ell = 1, \dots, p\}$, represent connections across layers at the same lattice point and control conditional dependence between variables j and ℓ at the i th lattice point (middle frame in Figure 1). Finally, the coefficients in the third summation, $\{b_{ijk\ell}; i = 1, \dots, n, k \in N_i, j \neq \ell = 1, \dots, p\}$, represent connections between locations across layers for different variables and control conditional cross-spatial dependence (right

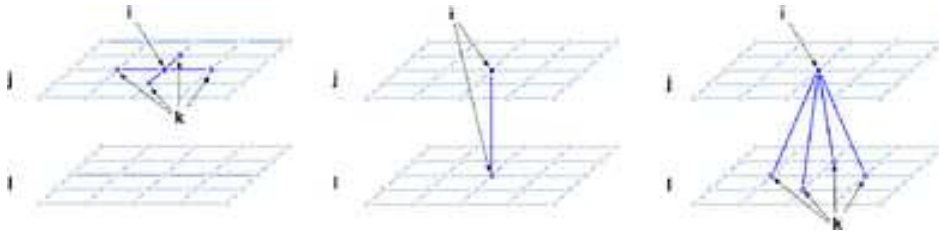


FIG. 1. Examples of different types of neighborhoods. The left frame shows a within-variable spatial neighborhood, while the middle frame shows a within-location neighborhood. The right frame demonstrates the neighborhood associated with cross-variable connections. See also equation (2).

frame in Figure 1). Of course, all of these conditional-dependence parameters and the variances $\{\tau_{ij}^2\}$ must satisfy regularity conditions that yield a symmetric, positive-definite covariance matrix for the joint distribution.

Some simplification of this basic structure is necessary to reduce the dimensionality of the parameter space. Ordering the data as $\mathbf{y} = [\mathbf{y}'_1, \dots, \mathbf{y}'_n]'$, where $\mathbf{y}_i = [y_{i1}, \dots, y_{ip}]'$ represents variables at the i th lattice point, we see from (1) that the joint distribution is Gaussian with mean given by $\boldsymbol{\mu} = [\boldsymbol{\mu}'_1, \dots, \boldsymbol{\mu}'_n]'$ and $\boldsymbol{\mu}_i = [\mu_{i1}, \dots, \mu_{ip}]'$, and with an $np \times np$ covariance matrix given by

$$(3) \quad \begin{bmatrix} \mathbf{A}_1 & \mathbf{B}_{12}\delta_{12} & \cdots & & \mathbf{B}_{1n}\delta_{1n} \\ \mathbf{B}_{21}\delta_{21} & \mathbf{A}_2 & & & \vdots \\ \vdots & & \ddots & & \\ & & & \mathbf{A}_{n-1} & \mathbf{B}_{n-1,n}\delta_{n-1,n} \\ \mathbf{B}_{n1}\delta_{n1} & \cdots & & \mathbf{B}_{n,n-1}\delta_{n,n-1} & \mathbf{A}_n \end{bmatrix}^{-1} \mathbf{T},$$

where $\delta_{ik} = 1$ if $k \in N_i$ and 0 otherwise. Each $p \times p$ block is given by

$$\mathbf{A}_i = \begin{bmatrix} 1 & & -b_{ijil} \\ & \ddots & \\ -b_{i\ell ij} & & 1 \end{bmatrix} \quad \text{or} \quad \mathbf{B}_{ik} = \begin{bmatrix} -b_{i1k1} & & -b_{ijk\ell} \\ & \ddots & \\ -b_{i\ell kj} & & -b_{ipkp} \end{bmatrix},$$

where $-b_{ijil}$ and $-b_{i\ell ij}$ are arbitrary off-diagonal elements of \mathbf{A}_i , and $-b_{ijk\ell}$ and $-b_{i\ell kj}$ are arbitrary off-diagonal elements of \mathbf{B}_{ik} . Finally, $\mathbf{T} = \text{diag}(\tau_{11}^2, \dots, \tau_{1p}^2, \dots, \tau_{n1}^2, \dots, \tau_{np}^2)$.

In general, we shall assume that the neighborhood structure is symmetric; that is, if the k th lattice point is a neighbor of the i th lattice point, then the i th lattice point is a neighbor of the k th ($\delta_{ik} = \delta_{ki}$). We shall also assume that $\tau_{ij}^2 = \tau_j^2$ for all j , implying a separate variance for each variable that does not vary with location. Hence, $\mathbf{T} = \mathbf{I}_n \otimes \text{diag}(\boldsymbol{\tau})$, where $\boldsymbol{\tau} = [\tau_1^2, \dots, \tau_p^2]'$. Further assumptions on the spatial-dependence parameters are made to reduce the dimensionality of the parameter space.

To address symmetry and variance homogeneity across location, it suffices to examine the components of specific blocks in the inverse of the covariance matrix given by (3). First, the *diagonal blocks* are given by

$$\text{diag}(\boldsymbol{\tau})^{-1} \mathbf{A}_i = \begin{bmatrix} 1/\tau_1^2 & & -b_{ijil}/\tau_j^2 \\ & \ddots & \\ -b_{i\ell ij}/\tau_\ell^2 & & 1/\tau_p^2 \end{bmatrix}.$$

By symmetry, the corresponding off-diagonal elements should be equal; that is, $b_{ijil}/\tau_j^2 = b_{i\ell ij}/\tau_\ell^2$. Setting $b_{ijil} = \rho_{j\ell}\tau_j/\tau_\ell$, with $\rho_{\ell j} = \rho_{j\ell}$, is one way of achieving the desired result. Then, $\text{diag}(\boldsymbol{\tau})^{-1} \mathbf{A}_i = \text{diag}(\boldsymbol{\tau})^{-1/2} \mathbf{A} \text{diag}(\boldsymbol{\tau})^{-1/2}$,

where

$$\mathbf{A} = \begin{bmatrix} 1 & & -\rho_{j\ell} \\ & \ddots & \\ -\rho_{j\ell} & & 1 \end{bmatrix}.$$

For the *off-diagonal blocks*, symmetry demands that

$$(4) \quad \text{diag}(\boldsymbol{\tau})^{-1} \mathbf{B}_{ik} = [\text{diag}(\boldsymbol{\tau})^{-1} \mathbf{B}_{ki}]'.$$

Assuming $i > k$, the left-hand side of (4) is given by

$$\text{diag}(\boldsymbol{\tau})^{-1} \mathbf{B}_{ik} = \begin{bmatrix} -b_{i1k1}/\tau_1^2 & & -b_{ijk\ell}/\tau_j^2 \\ & \ddots & \\ -b_{i\ell kj}/\tau_\ell^2 & & -b_{ipkp}/\tau_p^2 \end{bmatrix}.$$

Likewise, setting $b_{ijkl} = \phi_{jl}\tau_j/\tau_\ell$ gives $\text{diag}(\boldsymbol{\tau})^{-1} \mathbf{B}_{ik} = \text{diag}(\boldsymbol{\tau})^{-1/2} \mathbf{B} \times \text{diag}(\boldsymbol{\tau})^{-1/2}$, where

$$\mathbf{B} = \begin{bmatrix} -\phi_{11} & & -\phi_{j\ell} \\ & \ddots & \\ -\phi_{\ell j} & & -\phi_{pp} \end{bmatrix},$$

and, for $i > k$, set $\text{diag}(\boldsymbol{\tau})^{-1} \mathbf{B}_{ki} = \text{diag}(\boldsymbol{\tau})^{-1/2} \mathbf{B}' \text{diag}(\boldsymbol{\tau})^{-1/2}$ to satisfy (4).

The covariance in (3) then simplifies to

$$(5) \quad [\mathbf{I}_n \otimes \boldsymbol{\tau}^{1/2}] \left[\mathbf{I}_n \otimes \mathbf{A} - \begin{bmatrix} \mathbf{0} & & \mathbf{B}\delta_{ij} \\ & \ddots & \\ \mathbf{B}'\delta_{ij} & & \mathbf{0} \end{bmatrix} \right]^{-1} [\mathbf{I}_n \otimes \boldsymbol{\tau}^{1/2}],$$

where $\boldsymbol{\tau}^{1/2} = [\tau_1, \dots, \tau_p]'$ and where $\{\delta_{ij}\}$ are indicator functions for neighborhood dependence.

The specifications above simply ensure symmetry and reduce the number of parameters that must be estimated. Of course, the collection of spatial-dependence parameters, $\{\rho_{j\ell}\}$ and $\{\phi_{j\ell}\}$, must be chosen to ensure that (5) is a positive-definite covariance matrix. The final model has $p(p-1)/2$ within-location dependence parameters, $\{\rho_{j\ell}\}$, and p^2 between-location spatial-dependence parameters, $\{\phi_{j\ell}\}$, in addition to the p variance parameters, $\{\tau_j^2\}$, and any parameters that are used to define the means.

4. A hierarchical model for an RCM experiment. Let the n -dimensional vector \mathbf{y}_{rj} denote the output of an RCM, in particular, the r th ensemble member for the j th variable. In this work we focus solely on simple ensembles; that is, each member of the ensemble represents a perturbation of initial conditions for a single model. Potential extensions of this basic framework for perturbed physics or multi-model ensembles is discussed in Section 6.

At the first level of the hierarchy, the data model assumes that the vectors $\mathbf{y}_{rj}, r = 1, \dots, m, j = 1, \dots, p$, are independent with

$$(6) \quad \mathbf{y}_{rj} | \boldsymbol{\alpha}_j, \boldsymbol{\beta}_{rj}, \mathbf{h}_{rj}, \sigma_j^2 \sim \mathcal{N}(\mathbf{X}_1 \boldsymbol{\alpha}_j + \mathbf{X}_2 \boldsymbol{\beta}_{rj} + \mathbf{h}_{rj}, \sigma_j^2 \mathbf{I}),$$

where m indicates the number of ensemble members. In the mean structure, we allow for fixed effects common to all ensemble members within the j th variable ($\mathbf{X}_1 \boldsymbol{\alpha}_j$) and random effects specific to the r th ensemble member within the j th variable ($\mathbf{X}_2 \boldsymbol{\beta}_{rj}$). Spatial random effects are included through \mathbf{h}_{rj} , and σ_j represents a variable-specific variance.

The process model has two parts. First, the vectors $[\boldsymbol{\beta}'_{r1}, \dots, \boldsymbol{\beta}'_{rp}]', r = 1, \dots, m$, are assumed to be independent with

$$(7) \quad \begin{pmatrix} \boldsymbol{\beta}_{r1} \\ \vdots \\ \boldsymbol{\beta}_{rp} \end{pmatrix} \left| \begin{pmatrix} \boldsymbol{\beta}_1 \\ \vdots \\ \boldsymbol{\beta}_p \end{pmatrix} \right., \quad \boldsymbol{\Sigma}_b \sim \mathcal{N} \left(\begin{pmatrix} \boldsymbol{\beta}_1 \\ \vdots \\ \boldsymbol{\beta}_p \end{pmatrix}, \boldsymbol{\Sigma}_b \right),$$

where $\boldsymbol{\Sigma}_b$ is a $pq \times pq$ covariance matrix with q the number of columns of \mathbf{X}_2 . Second, the vectors $[\mathbf{h}'_{r1}, \dots, \mathbf{h}'_{rp}]', r = 1, \dots, m$, are assumed to be independent with

$$(8) \quad \begin{pmatrix} \mathbf{h}_{r1} \\ \vdots \\ \mathbf{h}_{rp} \end{pmatrix} \left| \begin{pmatrix} \mathbf{h}_1 \\ \vdots \\ \mathbf{h}_p \end{pmatrix} \right., \quad \{\tau_j^2\}, \{\rho_{j\ell}\}, \{\phi_{j\ell}\} \sim \mathcal{N} \left(\begin{pmatrix} \mathbf{h}_1 \\ \vdots \\ \mathbf{h}_p \end{pmatrix}, \mathbf{V}(\{\tau_j^2\}, \{\rho_{j\ell}\}, \{\phi_{j\ell}\}) \right).$$

(Note that the vectors $[\boldsymbol{\beta}'_{r1}, \dots, \boldsymbol{\beta}'_{rp}]'$ and $[\mathbf{h}'_{r1}, \dots, \mathbf{h}'_{rp}]'$ are assumed to be independent as well.) The first part, (7), focuses on linking the random regression coefficients specific to each ensemble member, while the second part, (8), imposes a multivariate structure on the spatial random effects. The covariance matrix \mathbf{V} takes its form from the multivariate Markov random field in (5).

The final level of the hierarchy assumes prior distributions on $\{\sigma_j\}$, $\{\boldsymbol{\alpha}_j\}$, $\{\boldsymbol{\beta}_j\}$, $\{\mathbf{h}_j\}$, $\boldsymbol{\Sigma}_b$, and the parameters of the spatial covariance, namely, $\{\tau_j^2\}$, $\{\rho_{j\ell}\}$, and $\{\phi_{j\ell}\}$. Typically, these priors will be vague or noninformative as well as independent. In addition, the prior distribution on $\{\rho_{j\ell}\}$ and $\{\phi_{j\ell}\}$ must ensure that the resulting covariance matrix is positive-definite.

From Bayes' theorem, the posterior distribution for the three-level hierarchical model is given by

$$\begin{aligned} & P(\{\boldsymbol{\beta}_{rj}\}, \{\mathbf{h}_{rj}\}, \{\boldsymbol{\alpha}_j\}, \{\sigma_j\}, \{\boldsymbol{\beta}_j\}, \boldsymbol{\Sigma}_b, \{\mathbf{h}_j\}, \{\tau_j^2\}, \{\rho_{j\ell}\}, \{\phi_{j\ell}\} | \mathbf{Y}) \\ & \propto P(\mathbf{Y} | \{\boldsymbol{\alpha}_j\}, \{\boldsymbol{\beta}_{rj}\}, \{\mathbf{h}_{rj}\}, \{\sigma_j\}) \end{aligned}$$

$$\begin{aligned} & \times P(\{\beta_{rj}\}|\{\beta_j\}, \Sigma_b)P(\{\mathbf{h}_{rj}\}|\{\mathbf{h}_j\}, \{\tau_j^2\}, \{\rho_{j\ell}\}, \{\phi_{j\ell}\}) \\ & \times P(\{\alpha_j\})P(\{\sigma_j\})P(\{\beta_j\})P(\Sigma_b)P(\{\mathbf{h}_j\})P(\{\tau_j^2\})P(\{\rho_{j\ell}\}, \{\phi_{j\ell}\}). \end{aligned}$$

It is clear that there is no closed-form solution for the posterior, and here Markov chain Monte Carlo (MCMC) [e.g., Gilks, Richardson and Spiegelhalter (1996)] is used to simulate realizations from the posterior distribution. In particular, we implement a Gibbs sampler [Geman and Geman (1984); Gelfand and Smith (1990); Gelfand et al. (1990)], incorporating Metropolis–Hastings steps [Metropolis et al. (1953); Hastings (1970)] where necessary.

One benefit of a MRF is that the specification involves the precision or inverse covariance matrix, and this matrix is typically sparse; that is, many of the elements of the matrix are zero. Methods for storing and manipulating such matrices have been widely established [e.g., Davis (2006)], and there is great potential for computational efficiency associated with sparse-matrix methods. There are now several sparse-matrix packages in the R statistical computing environment [R Development Core Team (2007)]. However, the **spam** package [Furrer (2008)] has functionality that is well suited for implementing MCMC with a MRF model. For example, the sparse Cholesky decomposition is one of the most important computational devices used when implementing the Gibbs sampler for MRF models such as those developed in this work. A typical sparse Cholesky decomposition involves three steps: (1) reorganizing the matrix by permuting the rows/columns to achieve a pattern of sparsity that is more efficient for the sparse Cholesky algorithm; (2) a symbolic step that identifies the pattern of sparsity in the matrix; and (3) the numerical computation. The first two steps do not change when manipulating matrices repeatedly with the same patterns of sparsity (as is the case here). The **spam** package allows one to achieve even greater computational efficiency by not repeating these steps during the course of the MCMC. For more details on this and other computational benefits gained from incorporating sparse-matrix methods in such applications, see Furrer and Sain (2010).

5. The RCM experiment. Leung et al. (2004) describe an RCM experiment using the NCAR/DOE Parallel Climate Model to drive the NCAR/Penn State Mesoscale Model (MM5) as an RCM. The experiment produced a control run from 1995–2015 and three future runs (ensemble members) from 2040–2060. The domain consisted of the western United States and part of western Canada, and the model used a “business as usual” climate scenario incorporating a 1% annual increase in the amount of greenhouse gases.

The $n = 44 \times 56 = 2464$ grid boxes form a regular lattice. Since a long-run average of weather is one way to quantify climate, typical summaries of climate model runs include seasonal averages of temperature and precipitation with the length of the integration often determined by computational

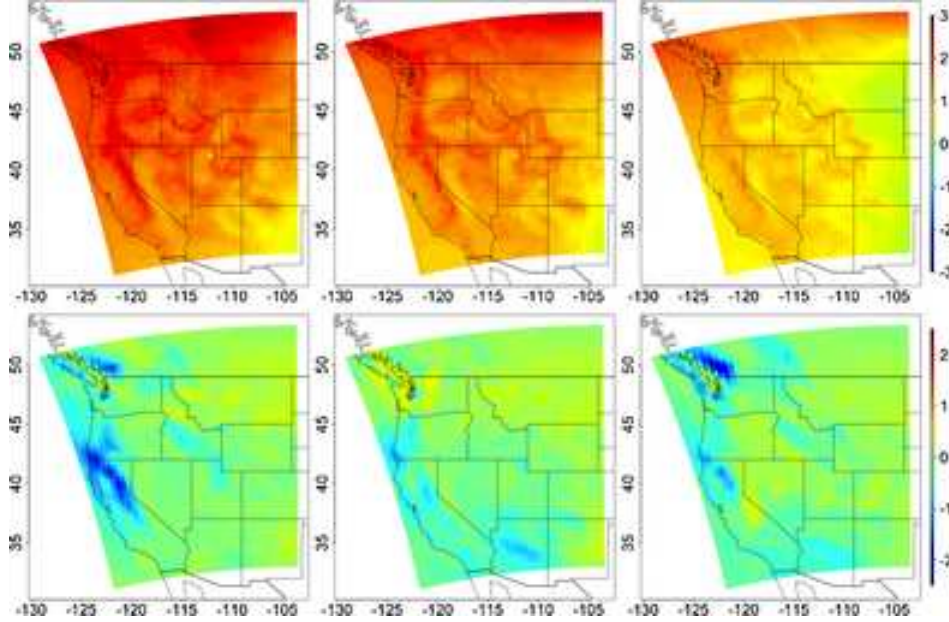


FIG. 2. Top row shows the three differences in winter midpoint temperature ($^{\circ}$ K), while the bottom row shows the three differences in total precipitation (inches).

considerations. Hence, twenty-year winter (December, January, and February) average temperature and average total precipitation were computed for each grid box and for each of the control and the three future runs. Differences between the future and the control were calculated, yielding change-in-temperature and change-in-total-precipitation variables. Hence, there are $p = 2$ variables and $m = 3$ ensemble members, giving six fields to be analyzed. These spatial fields for the winter season are shown in Figure 2. A second, separate analysis with the same structure is also presented for the twenty-year summer (June, July, and August) change in temperature and change in total precipitation.

We note that the statistical models outlined in this work focus on using Gaussian assumptions for average precipitation from an RCM. (An assumption that was verified through exploratory analysis.) However, other models for precipitation, in particular, for extreme precipitation, are possible; see, for example, Sansó and Guenni (2004), Schliep et al. (2010), and Cooley and Sain (2010).

5.1. Model specification. We now outline some specifics about the statistical-model specification. Consider the data model (6). After some exploratory analysis, (scaled) latitude, longitude, and elevation were used as covariates in the common regression component ($\mathbf{X}_1\boldsymbol{\alpha}_j$), to which a random intercept

across ensemble members was added ($\mathbf{X}_2\boldsymbol{\beta}_{rj}$). The prior covariance matrix for the random intercept in (7) was also simplified to $\boldsymbol{\Sigma}_b = \sigma_b^2 \mathbf{I}_p$.

The prior distribution for the variance parameters, $\{\sigma_j^2\}$ and σ_b^2 , were taken to be noninformative; that is, they were assumed to follow the prior distribution $P(\sigma^2) \propto 1/\sigma^2$, independently. The prior distributions for the regression parameters, $\{\boldsymbol{\alpha}_j\}$ and $\{\boldsymbol{\beta}_j\}$, were taken to be mean-zero Gaussian distributions with covariance matrices proportional to the identity and with large variances (i.e., $\sigma_\alpha^2 = 10$ and $\sigma_\beta^2 = 100$). The prior distributions for $\{\mathbf{h}_j\}$ were also taken to be mean-zero Gaussian distributions with covariance matrices proportional to the identity and with large variances (i.e., $\sigma_{\mathbf{h}_j}^2 = 10$).

Finally, the prior specification for the joint distribution of ρ , ϕ_{11} , ϕ_{22} , ϕ_{12} , and ϕ_{21} was taken to be uniform over the range of values that yield a positive-definite covariance matrix. This region was identified using rejection sampling based on a sparse Cholesky decomposition. A simple simulation study of a univariate Markov random field's spatial dependence parameter (not reported here) suggested that concentrating priors on a subregion of the parameter space leads to a biased estimate when the true parameter lies outside this region. While this may not seem surprising, the lesson learned is that there has to be a good reason to choose nonuniform priors for these bounded spatial-dependence parameters.

5.2. Results for the winter season. Posterior distributions were obtained using MCMC algorithms, and considerable care was taken to ensure the convergence of the parameters in the MCMC. This is especially true with respect to the conditional-dependence parameters, where our experience has shown that straightforward approaches can lead to disappointing performance (i.e., very slow mixing and convergence). Ten chains were run, each with random starting values; the starting values for the conditional-dependence parameters were chosen uniformly across the space of values that yield a positive-definite covariance matrix.

A Gibbs sampler was implemented that involved three distinct regimes. In the first regime (2500 iterations), each of the conditional-dependence parameters was updated one at a time using a Metropolis–Hastings algorithm. Gaussian proposal distributions were used, with periodic updates of the proposal variance to achieve an approximate 20% acceptance rate. In the second regime (the next 10,000 iterations), ρ , ϕ_{12} , and ϕ_{21} were updated simultaneously using a Metropolis–Hastings algorithm with a multivariate Gaussian proposal distribution. Again, the proposal covariance matrix was updated periodically to achieve an approximate 20% acceptance rate. Other conditional-dependence parameters were still updated using a univariate Metropolis–Hastings algorithm. Finally, in the third regime (the

last 10,000 iterations), ρ , ϕ_{12} , and ϕ_{21} were again updated simultaneously, but no further updates of the proposal distribution were made. Convergence of the posterior distributions of the parameters in the MCMC was monitored using both graphical and numerical methods [e.g., Gelman (1996)]. Posterior distributions were then estimated by sampling from the third regime.

Of particular interest are the conditional-dependence parameters, since these control the nature and degree of the spatial correlation in the model. Figure 3 shows scatterplots and kernel estimates of the distribution of ϕ_{11} (temperature) and ϕ_{22} (precipitation), the parameters that control the conditional dependence between lattice points within a layer (Figure 1, left panel). The distributions show that there is considerable (conditional) spatial dependence within each variable, as the distributions tend to be concentrated near the positive boundary of possible values for ϕ_{11} and ϕ_{22} . There is evidence of a slightly stronger dependence for temperature (ϕ_{11}).

Trace plots and other diagnostics for ρ , ϕ_{12} , and ϕ_{21} suggest convergence after about 10,000 iterations, which corresponds to the end of the second sampling regime. These three parameters control the dependence structure across variables; ρ summarizes the within-location dependence (Figure 1, middle panel) and ϕ_{12} , ϕ_{21} summarize the cross-variable dependence (Figure 1, right panel). The estimated posterior mean and posterior standard deviation for ρ is -0.12 and 0.014 , respectively. A negative value for ρ suggests that an increasing temperature is (conditionally) associated with a decreasing total precipitation.

Figure 4 highlights the distribution of ϕ_{12} and ϕ_{21} . The strong correlation between these two conditional cross-correlation parameters is clearly shown

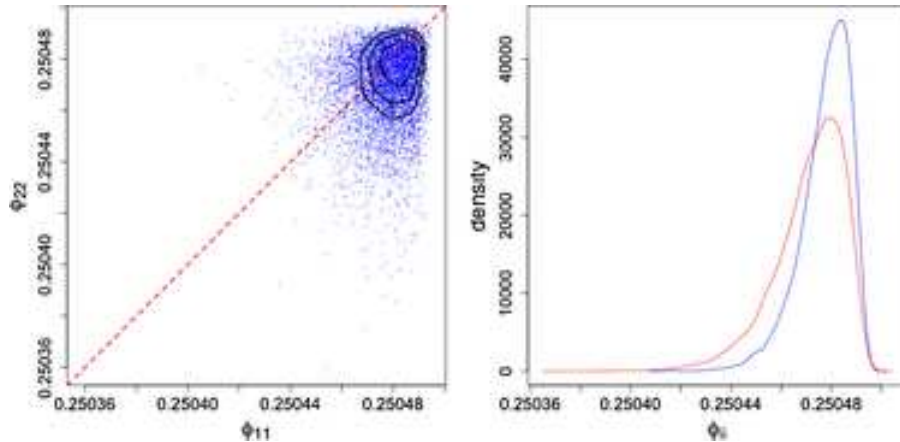


FIG. 3. Left frame shows scatterplot of a random sample of 10,000 values of ϕ_{11} and ϕ_{22} (1000 from each of the 10 chains). Contours represent approximate 25, 50, and 75% contours of a kernel density estimate. Right frame shows kernel density estimates of the marginals for ϕ_{11} (blue) and ϕ_{22} (red).

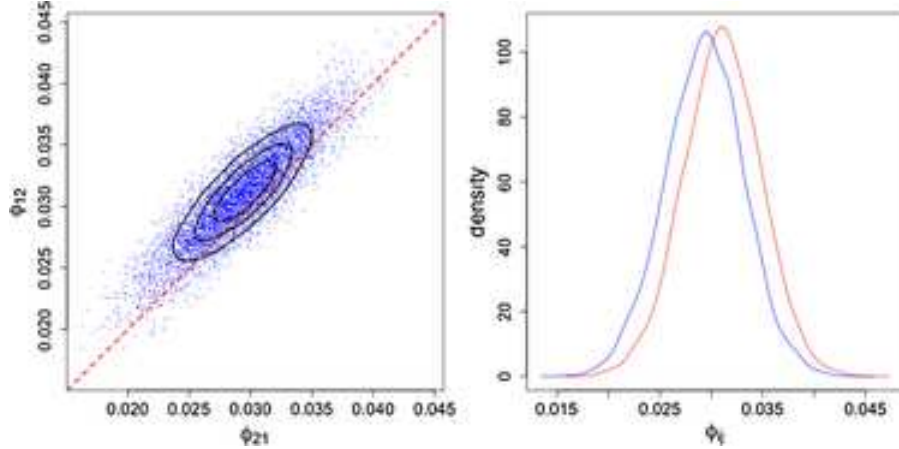


FIG. 4. Left frame shows scatterplot of a random sample of 10,000 values of ϕ_{12} and ϕ_{21} (1000 from each of the 10 chains). Contours represent approximate 25, 50, and 75% contours of a kernel density estimate. Right frame shows kernel density estimates of the marginals for ϕ_{12} (red) and ϕ_{21} (blue).

in the left frame of Figure 4. However, there is another feature of note: There is compelling evidence of asymmetry in the strength of these two parameters, with roughly 85% of the sampled points being above the line $y = x$. This suggests that there is higher conditional dependence between temperature values and neighboring total precipitation than there is conditional dependence between total precipitation values and neighboring temperature values. Almost all of the published models of multivariate MRFs assume $\phi_{12} = \phi_{21}$, something we have argued previously as being overly restrictive [Sain and Cressie (2007)]. Our posterior inference shows the inappropriateness of such an assumption in this case.

Figure 5 shows posterior means for the fixed regression components (left column), the spatial random effects (middle column), and their sum (right column), for the change in winter average temperature (top row) and the change in total winter precipitation (bottom row). The fixed effects show a clear latitudinal effect as well as an east-to-west gradient. For precipitation, there is a more dominant east-to-west gradient. The spatial random effects for the change in temperature seem to follow the features of the topography, and are, in general, of smaller magnitude than the fixed effects. The spatial effects for the change in total precipitation also follow the features of the topography, but there are additional strong local features, for example, in northern California. In contrast to temperature, the spatial effects for total precipitation are larger relative to the fixed effects.

The sum of the fixed effects and the spatial random effects for temperature shows a consistent pattern of winter warming on average throughout the

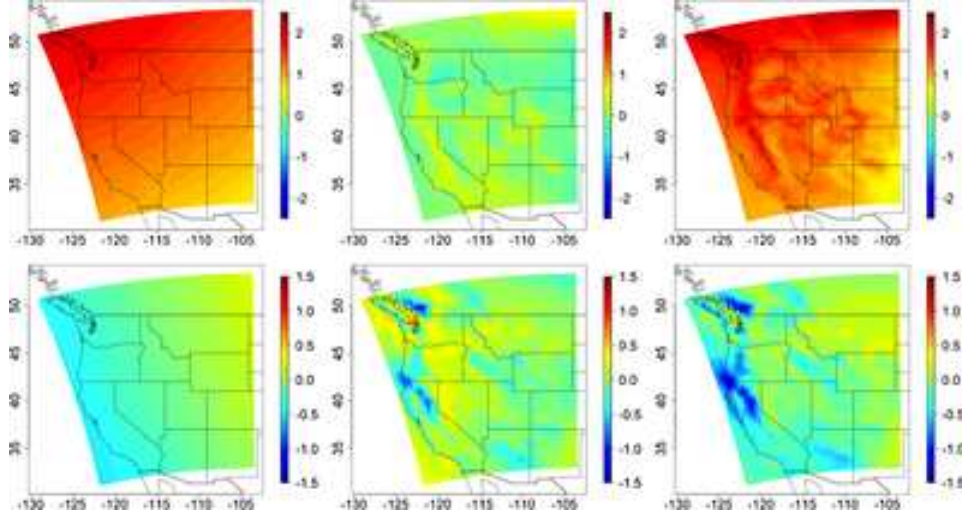


FIG. 5. *Posterior means for the regression (left), the spatial effect (middle), and the sum (right) for the winter season. The top row represents the change in midpoint temperature ($^{\circ}K$), while the bottom represents the change in total precipitation (inches).*

west, while the sum for total precipitation shows patterns that are much more localized. The most dominant signal for total precipitation is indicated by the regions of sharp decline in winter precipitation in northern California and the Pacific northwest.

To aid in the identification of areas that might be at most risk for change, as projected by this regional-climate-model experiment, Figure 6 shows the result of a hierarchical clustering based on the posterior distribution of the mean change in temperature and total precipitation for each grid. One thousand samples were drawn from the posterior distribution of the mean change in temperature and total precipitation for each grid box. The distance metric used to join clusters was a symmetrized version of the Kullback–Leibler distance, which was based on the assumption of bivariate normality within each cluster. Hence, the clustering focuses not only on the mean of the posterior distribution but also includes information about the covariance structure of the changes for each grid box. The scatterplot in the left frame of Figure 6 shows the posterior means for each grid box with each cluster indicated by different colors. The scatterplot shows the considerable structure that the clustering is able to distinguish. The spatial pattern of the clusters is also shown in Figure 6 (right frame). The dark red areas, for example, highlight a region associated with a strong increase in temperature and a strong decrease in total precipitation.

It is also useful to consider other measures of uncertainty. Fields of standard deviations are one approach, but when considering the multivariate

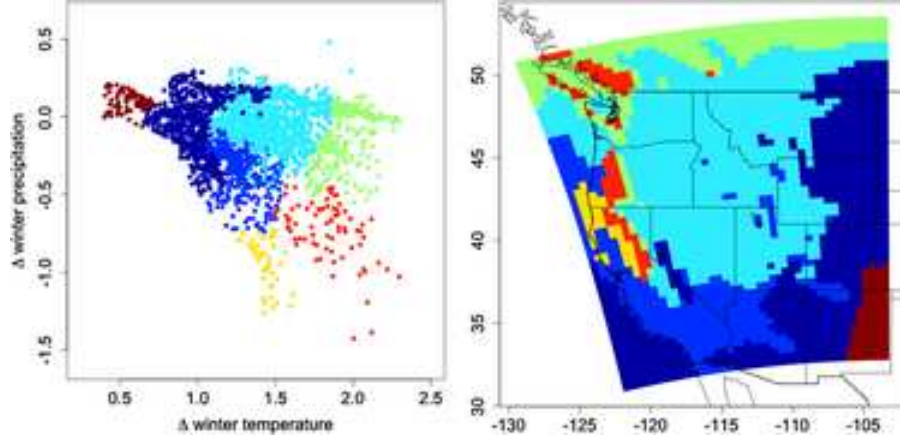


FIG. 6. Results from clustering the posterior means of the change in temperature and the change in total precipitation for the winter season. The left frame is a scatterplot with the clusters indicated through different colors. The right frame shows the clusters spatially.

nature of this model output other possibilities arise. For example, Figure 7 shows estimated pointwise probabilities of an increase in temperature (left frame), a decrease in total precipitation (middle frame), and a simultaneous increase in temperature and decrease in total precipitation (right frame) based on a sampling of the posterior distribution of the joint spatial fields. In general, temperature is increasing on average across the entire domain; hence, these probabilities are based on increases larger than the median computed from all the samples across all the grid boxes. Likewise, the decrease in total precipitation was based on decreases larger than the median across all the samples from all the grid boxes. Again, on the basis of this model, we see evidence of widespread increase in temperature and decrease in total precipitation across the western U.S., with dominant features along the far western coast, the Pacific northwest, and isolated mountain regions.

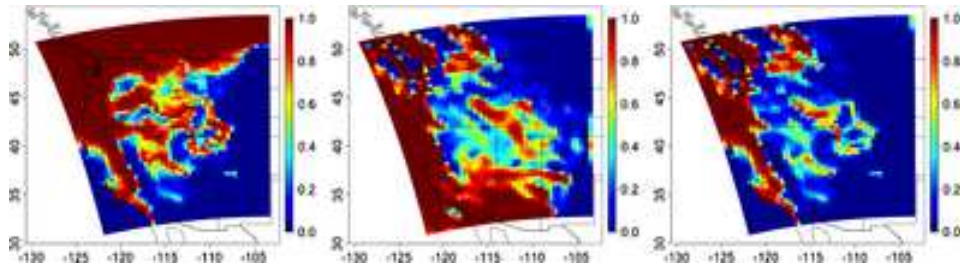


FIG. 7. Estimated pointwise probabilities for the winter season: increasing temperature (left), decreasing total precipitation (middle), and simultaneously increasing temperature and decreasing total precipitation (right).

Figure 8 shows an alternative representation of the joint distribution by considering conditional probabilities. The figure shows the probability of the decrease in total precipitation being in the top quartile, conditional on the increase in temperature being in the first quartile (top left), second quartile (top right), third quartile (bottom left) and fourth quartile (bottom right). As the temperature increase becomes more extreme, the largest decreases in total precipitation move from being focused in the southwest (and the California coast) to the Pacific northwest (and the California coast). Relative changes can also be considered and, although not shown here, the normalization minimizes the impact on the Pacific northwest (higher absolute changes in total precipitation, but also higher total precipitation values in general).

Finally, Figure 9 shows approximate 95% contours for the joint change in (average winter) temperature and total precipitation, but focuses on the grid boxes that represent the five consolidated metropolitan statistical areas (as defined by the U.S. Census) that are included in the domain. Figure 9 suggests that the projection for the Denver area includes average increases in

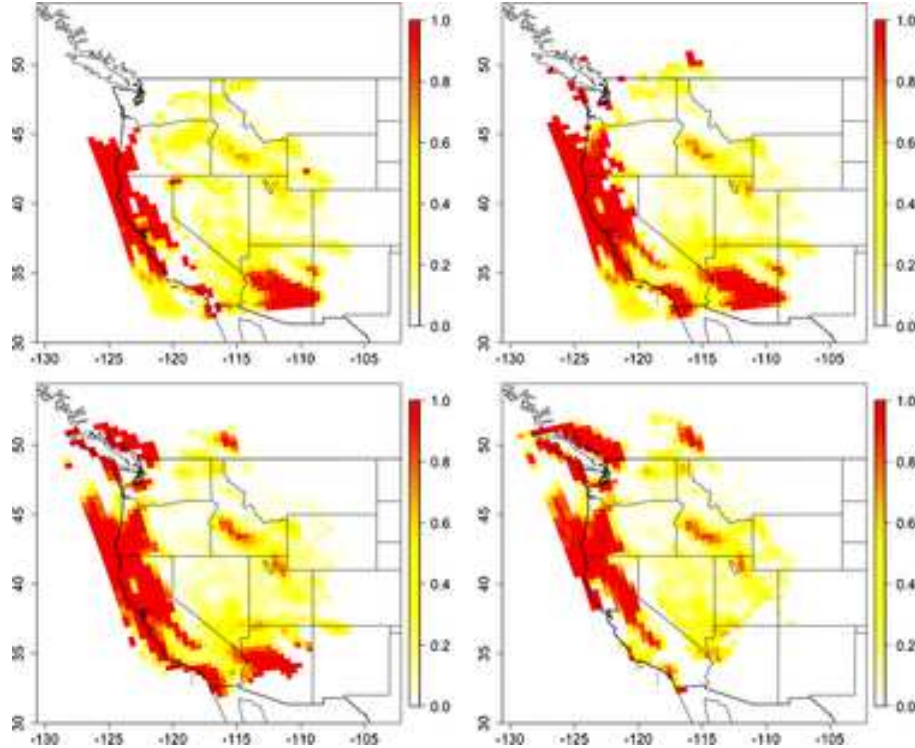


FIG. 8. Probability of a large decrease in winter total precipitation, conditional on the increase in temperature falling in the first quartile (top left), second quartile (top right), third quartile (bottom left), and fourth quartile (bottom right).

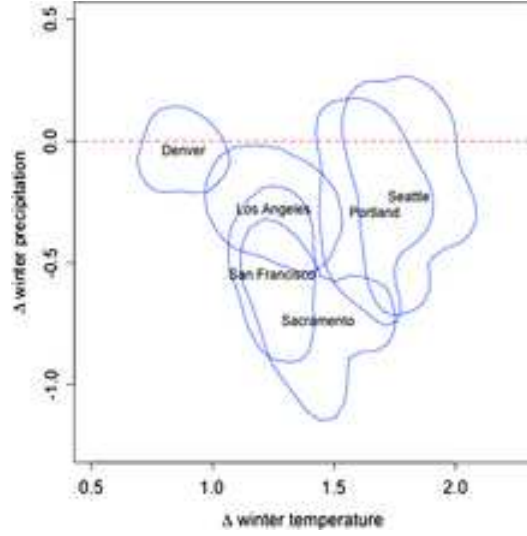


FIG. 9. *Approximate 95% contours for the average change in winter temperature and precipitation for the grid boxes associated with the five consolidated metropolitan statistical areas in the domain.*

temperature of just under 1°K with minimal average decreases in total precipitation. The contour for the Sacramento area, on the other hand, suggests much larger average increases in temperature and average decreases in total precipitation. We believe that such plots, summarizing the joint distribution based on the statistical model, will have great interest and application to scientists and decision-makers interested in the impacts of climate change.

5.3. Results for the summer season. A slightly different scheme was used for the Gibbs sampler for the analysis of the summer model output. The three-regime sampling was still used, although with twice as many iterations (20,000) in the second regime. In addition, all five conditional-dependence parameters (ρ , ϕ_{11} , ϕ_{22} , ϕ_{12} , and ϕ_{21}) were updated simultaneously. Convergence of parameters for the summer season was similar to that for the winter season, although somewhat slower, and the specifics of those results are not discussed here. Distributions of the posteriors for the parameters were similar, with the exception of the cross-dependence parameters, and, in particular, ρ (posterior mean of -0.41 for summer versus a posterior mean of -0.12 for winter), suggesting that the summer season has a much stronger and more negative correlation between the change in temperature and the change in total precipitation.

Figure 10 shows posterior means for the summer season with the same layout as in Figure 5 for the winter season. Now there appears to be a west-to-east gradient in the fixed effects for temperature, and, again, the spatial

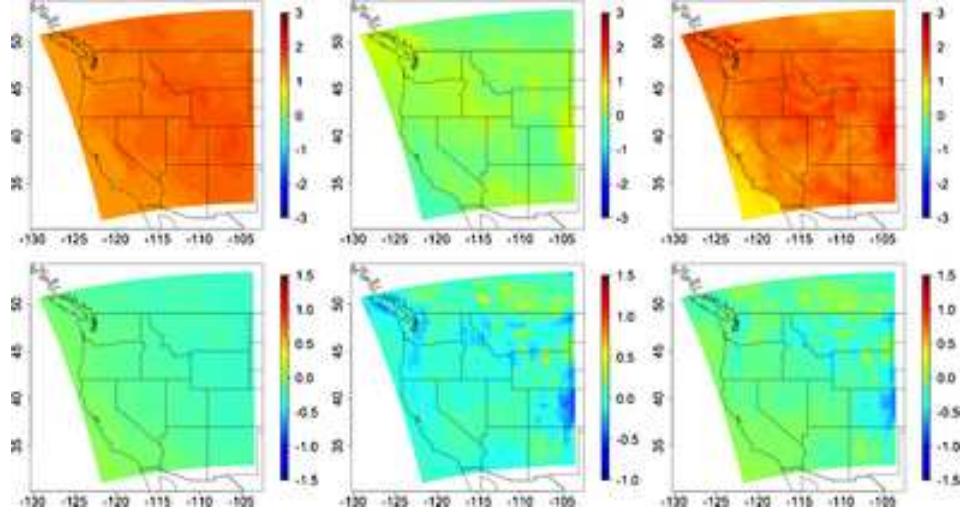


FIG. 10. *Posterior means for the regression (left), the spatial effect (middle), and the sum (right) for the summer season. The top row represents the change in midpoint temperature ($^{\circ}$ K), while the bottom represents the change in total precipitation (inches).*

random effects pick up more of the topography that is not accounted for in the fixed effects. Again, for temperature, the spatial random effects are of smaller magnitude than the fixed effects.

For total precipitation, there is also a west-to-east gradient in the fixed effects. The spatial random effects for the change in total precipitation also follow the features of the topography, but there are strong local features, now occurring in the eastern part of the domain. In comparison to temperature, the spatial random effects are larger relative to the fixed effects.

The sum of the fixed effects and the spatial random effects for temperature shows a consistent pattern of summer warming on average throughout the west, while the sum for total precipitation shows patterns that are much more localized, just as in the winter season. However, the most dominant features for total precipitation are the regions of decrease in summer total precipitation in the eastern part of the domain. Again, a hierarchical clustering was performed on samples from the posterior distribution for each grid box, which is summarized in Figure 11. There appears to be more widespread warming and decreasing total precipitation during the summer months, but the clustering again highlights structure in the joint distribution.

As in Figure 7, Figure 12 shows estimated pointwise probabilities of an increase in temperature (left frame), a decrease in total precipitation (middle frame), and a simultaneous increase in temperature and decrease in total precipitation (right frame) based on a sampling of the posterior distribution of the joint spatial fields. In general, summer warming and decreasing total

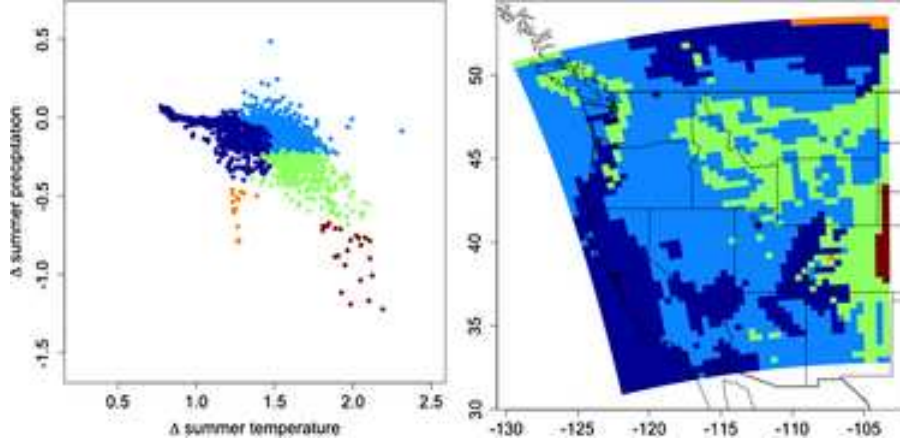


FIG. 11. Results from clustering the posterior means of the change in temperature and total precipitation for the summer season. The left frame is a scatterplot with the clusters indicated through different colors. The right frame shows the clusters spatially.

precipitation is widespread, even more so than in the winter season, and focused more on the eastern side of the domain.

Figure 13 is constructed similarly to Figure 8. However, the stronger negative correlation in the summer season is more apparent in the figure. The large decreases in total precipitation are strongest in the eastern portion of the domain, but decrease dramatically when we condition on larger increases in temperature.

Finally, Figure 14 shows approximate 95% contours for the joint change in (average summer) temperature and total precipitation. In this case, the roles are reversed from the winter season. The contours suggest larger increases in temperature and decreases in total precipitation (on average) for the Denver area, while the contour for the Sacramento area suggests more modest changes on both variables.

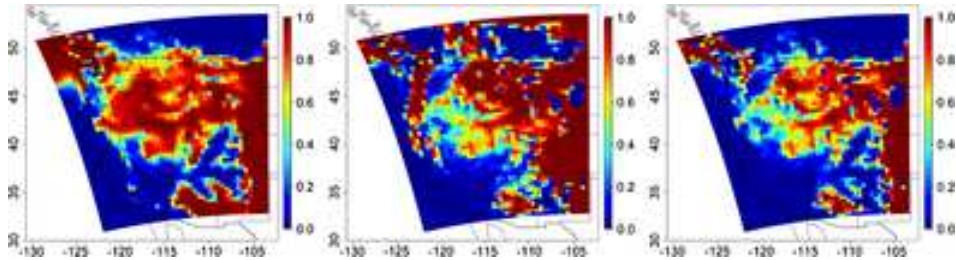


FIG. 12. Estimated pointwise probabilities for the summer season: increasing temperature (left), decreasing total precipitation (middle), and simultaneously increasing temperature and decreasing total precipitation (right).

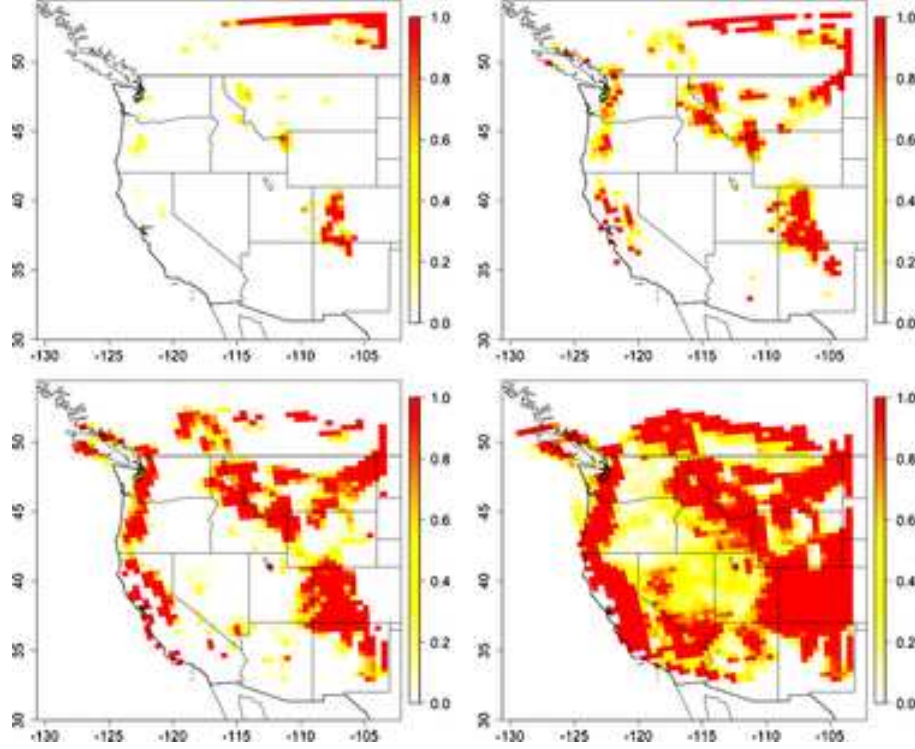


FIG. 13. *Probability of a large decrease in summer total precipitation, conditional on the increase in temperature falling in the first quartile (top left), second quartile (top right), third quartile (bottom left), and fourth quartile (bottom right).*

6. Concluding remarks. Climate models have become an important tool in the study of climate and climate change. Ensemble experiments of climate-model output, be they comprised of perturbed initial conditions, perturbed physics, or multiple models, have also become important in studying and quantifying the uncertainty in climate-model output. However, there are typically only a limited number of runs that can be produced due to the time and expense of running these models, even on modern supercomputers. Hence, statistical methods become necessary to quantify the distribution and the breadth of variation in the model output.

With this idea in mind, we have introduced a hierarchical spatial statistical model designed primarily for the analysis of regional-climate-model output on the basis of a simple ensemble (perturbed initial conditions). This model is multivariate and has the capacity to simultaneously characterize multiple model outputs, for example, the average change in temperature and the average change in total precipitation. While analysis of the individual model outputs might yield estimates of marginal distributions, the strong

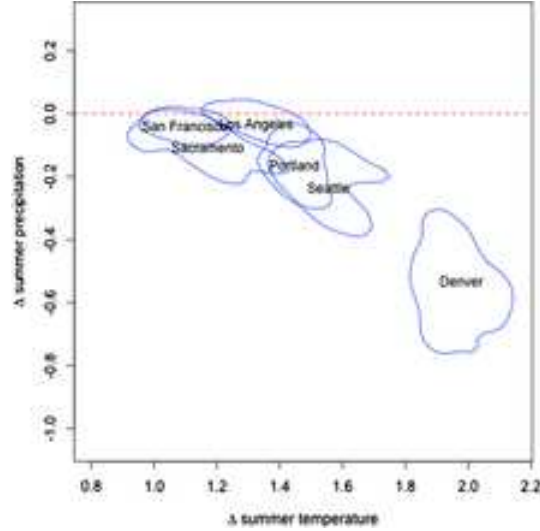


FIG. 14. *Approximate 95% contours for the average change in summer temperature and precipitation for the grid boxes associated with the five consolidated metropolitan statistical areas in the domain.*

correlations across variables, such as those uncovered here in the analysis of the average changes in summer temperature and the average changes in total precipitation, make a multivariate analysis crucial for joint inference.

The statistical model also captures the spatial variation in the model output through a novel implementation of a multivariate MRF. In addition to the computational benefits arising from using models based on an MRF, this formulation of a multivariate MRF has a great deal of flexibility in modeling the conditional-dependence structure and is easily extendable. For example, more complex neighborhood structures can be considered [e.g., Sain, Furrer and Cressie (2007)], and it is not difficult to conceptualize how one might even consider modeling the joint distribution of multiple variables that are on different lattices. Connections to graphical models [e.g., Whittaker (1990)] could lead to further insights into modeling and parameter estimation. The computational impact and practical utility of considering additional variables (increasing p) or additional ensemble members (increasing m) is also of interest, although, at least in climate model research, these are highly dependent on the application and the computational demands associated with running climate models on high-performance computers. These are issues that we are currently considering. We also note that this work adds to a growing collection of research involving the study and modeling of asymmetric cross-dependence structures for multivariate spatial data, including, for example, Jin, Carlin and Banerjee (2005) and Sain and Cressie (2007).

in the case of MRFs and Royle and Berliner (1999), Ver Hoef, Cressie and Barry (2004), and Apanasovich and Genton (2010) for geostatistical data.

There is great interest in more complex ensembles, such as perturbed-physics experiments and multi-model ensembles. We have not considered such ensembles here, as we have focused on the multivariate aspect of the analysis of simple ensembles of regional-climate-model output. However, the model presented here could be extended to consider such ensembles by straightforward modifications to the process model, in particular, equation (8) could be modified to allow for a spatial meta-analysis component [e.g., Kang, Cressie and Sain (2010)] or through a functional analysis of variance similar to that of Kaufman and Sain (2010). Aside from the obvious computational challenges to simply fitting such models in the multivariate setting, there is, of course, more work needed to quantify the variation associated with different model physics or different models.

It is important to note that any conclusions taken from an analysis such as the one considered here are conditional on the assumptions implicit to the particular climate model or models used to generate the output fields. Whether global or regional in nature, climate models are typically constructed to reproduce certain features in the current climate, and analyzing differences as we did should minimize the impact of any biases in the climate models (although this approach might be viewed with some healthy skepticism, as it is not clear that the biases in current runs are going to be the same as biases in future runs). Observations may be included to help constrain the statistical model, at least with respect to current climate. However, it is still an open question how to include observational data sets for spatial analyses of RCMs of the sort done here. Station-level data does not have the same spatial and temporal coverage, and there are also numerous additional issues with using interpolated data products, including reanalysis data that represent a data assimilation using both station-level data and climate-model output.

In addition to the longitude, latitude, and elevation used in this analysis, predictors based on climatology (long-run means of temperature and precipitation) were considered [e.g., Furrer et al. (2007a, 2007b)], but these were ultimately ruled out as not being effective at predicting the changes in temperature and precipitation. However, work done by Tebaldi et al. (2005) offers an approach for combining model output and observations. We are currently considering how their approach may be useful for spatial analyses of RCM output.

Finally, there is also much interest, for example, from people examining the impacts of climate change, in combining model output in order to obtain improved projections of climate change or to span the variation across a climate-model experiment. Of course, with the more complex climate-model experiments, there is the issue of model-to-model correlations. We believe

that the inherent multivariate nature of this model provides an excellent starting place to consider such correlations.

Acknowledgments. We would like to express our appreciation to the referees and the editor for their helpful comments. The National Center for Atmospheric Research is sponsored by the National Science Foundation (NSF).

REFERENCES

- APANASOVICH, T. V. and GENTON, M. G. (2010). Cross-covariance functions for multivariate random fields based on latent dimensions. *Biometrika* **97** 15–30.
- BANERJEE, S., CARLIN, B. P. and GELFAND, A. E. (2004). *Hierarchical Modeling and Analysis for Spatial Data*. Chapman & Hall/CRC Press, Bacon Raton, FL.
- BERLINER, L. M. and KIM, Y. (2008). Bayesian design and analysis for superensemble based climate forecasting. *Journal of Climate* **21** 1891–1910.
- BESAG, J. E. (1974). Spatial interaction and the statistical analysis of lattice systems (with discussion). *J. Roy. Statist. Soc. Ser. B* **35** 192–236. [MR0373208](#)
- BILLHEIMER, D., CARDOSO, T., FREEMAN, E., GUTTORP, P., KO, H. and SILKEY, M. (1997). Natural variability of benthic species composition in the Delaware Bay. *Environ. Ecol. Stat.* **4** 95–115.
- CARLIN, B. P. and BANERJEE, S. (2003). Hierarchical multivariate CAR models for spatiotemporally correlated data. In *Bayesian Statistics* **7** 45–63. Oxford University Press, Oxford. [MR2003166](#)
- COOLEY, D. and SAIN, S. R. (2010). Spatial hierarchical modeling of precipitation extremes from a regional climate model. *J. Agric. Biol. Environ. Stat.* **15** 381–402.
- CRESSIE, N. A. C. (1993). *Statistics for Spatial Data*, rev. ed. Wiley, New York. [MR1239641](#)
- DANIELS, M. J., ZHOU, Z. and ZOU, H. (2006). Conditionally specified space–time models for multivariate processes. *J. Comput. Graph. Statist.* **15** 157–177. [MR2269367](#)
- DAVIS, T. A. (2006). *Direct Methods for Sparse Linear Systems*. SIAM, Philadelphia. [MR2270673](#)
- FURRER, R. (2008). spam: SPARse Matrix. R package version 0.14-1.
- FURRER, R. and SAIN, S. R. (2010). spam: A sparse matrix r package with emphasis on mcmc methods for Gaussian Markov random fields. *Journal of Statistical Software* **36** 1–25.
- FURRER, R., KNUTTI, R., SAIN, S. R., NYCHKA, D. and MEEHL, G. A. (2007a). Spatial patterns of probabilistic temperature change projections from a multivariate Bayesian analysis. *Geophysical Research Letters* **34** L06711. DOI: [10.1029/2006GL027754](#).
- FURRER, R., SAIN, S. R., NYCHKA, D. and MEEHL, G. A. (2007b). Multivariate Bayesian analysis of atmosphere–ocean general circulation models. *Environ. Ecol. Stat.* **14** 249–266. [MR2405329](#)
- GELFAND, A. E., HILLS, S. E., RACINE-POON, A. and SMITH, A. F. M. (1990). Illustration of Bayesian inference in normal data models using Gibbs sampling. *J. Amer. Statist. Assoc.* **85** 972–985.
- GELFAND, A. E. and SMITH, A. F. M. (1990). Sampling-based approaches to calculating marginal densities. *J. Amer. Statist. Assoc.* **85** 398–409. [MR1141740](#)
- GELFAND, A. E. and VOUNATSOU, P. (2003). Proper multivariate conditional autoregressive models for spatial data analysis. *Biostatistics* **4** 11–15.
- GELMAN, A. (1996). Inference and monitoring convergence. In *Markov Chain Monte Carlo in Practice* 131–144. Chapman & Hall, London.

- GEMAN, S. and GEMAN, D. (1984). Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **6** 721–741.
- GILKS, W. R., RICHARDSON, S. and SPIEGELHALTER, D. J. (1996). Introducing Markov chain Monte Carlo. In *Markov Chain Monte Carlo in Practice* 1–19. Chapman & Hall, London. [MR1397966](#)
- HASTINGS, W. K. (1970). Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* **57** 97–109.
- HOUGHTON, J. T., DING, Y., GRIGGS, D. J., NOGUER, M., VAN DER LINDEN, P. J., DAI, X., MASKELL, K. and JOHNSON, C. A., EDS. (2001). *Climate Change 2001: The Scientific Basis. Contribution of Working Group I to the Third Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge Univ. Press, Cambridge.
- JIN, X., BANERJEE, S. and CARLIN, B. P. (2007). Order-free coregionalized lattice models with application to multiple disease mapping. *J. Roy. Statist. Soc. Ser. B* **69** 817–838. [MR2368572](#)
- JIN, X., CARLIN, B. P. and BANERJEE, S. (2005). Generalized hierarchical multivariate CAR models for areal data. *Biometrics* **61** 950–961. [MR2216188](#)
- KANG, E. L., CRESSIE, N. and SAIN, S. (2010). Combining outputs from the NARCCAP regional climate models using a Bayesian hierarchical model. Technical Report 837, Dept. Statistics, Ohio State Univ., Columbus, OH.
- KAUFMAN, C. G. and SAIN, S. R. (2010). Bayesian functional ANOVA modeling using Gaussian process prior distributions. *Bayesian Anal.* **5** 123–150.
- KIM, H., SUN, D. and TSUTAKAWA, R. K. (2001). A bivariate Bayes method for improving the estimates of mortality rates with a twofold conditional autoregressive model. *J. Amer. Statist. Assoc.* **96** 1506–1521. [MR1946594](#)
- LEUNG, L., QIAN, Y., BIAN, X., WASHINGTON, W. M., HAN, J. and ROADS, J. O. (2004). Mid-century ensemble regional climate change scenarios for the western United States. *Climatic Change* **62** 75–113.
- MARDIA, K. V. (1988). Multidimensional multivariate Gaussian Markov random fields with applications to image processing. *J. Multivariate Anal.* **24** 265–284. [MR0926357](#)
- METROPOLIS, N., ROSENBLUTH, A. W., ROSENBLUTH, M. N., TELLER, A. H. and TELLER, E. (1953). Equations of state calculations by fast computing machines. *Journal of Chemical Physics* **21** 1087–1091.
- PETTITT, A. N., WEIR, I. S. and HART, A. G. (2002). A conditional autoregressive Gaussian process for irregularly spaced multivariate data with application to modeling large sets of binary data. *Stat. Comput.* **12** 353–367. [MR1951708](#)
- R Development Core Team (2007). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- ROYLE, A. M. and BERLINER, L. M. (1999). A hierarchical approach to multivariate spatial modeling and prediction. *J. Agric. Biol. Environ. Stat.* **4** 29–56. [MR1812239](#)
- RUE, H. and HELD, L. (2005). *Gaussian Markov Random Fields: Theory and Application*. Chapman & Hall/CRC Press, Boca Raton, FL. [MR2130347](#)
- SAIN, S. R. and CRESSIE, N. (2007). A spatial model for multivariate lattice data. *J. Econometrics* **140** 226–259. [MR2395923](#)
- SAIN, S. R., FURRER, R. and CRESSIE, N. (2007). Combining regional climate model output via a multivariate Markov random field model. In *Bulletin of the International Statistical Institute* **LXII** 1375–1382. Instituto Nacional de Estatística, Lisboa, Portugal.
- SANSÓ, B. and GUENNI, L. (2004). A Bayesian approach to compare observed rainfall data to deterministic simulations. *Environmetrics* **15** 597–612.

- SCHABENBERGER, O. and GOTWAY, C. A. (2005). *Statistical Methods for Spatial Data Analysis*. Chapman & Hall/CRC Press, Boca Raton, FL. [MR2134116](#)
- SCHLIEP, E. M., COOLEY, D., SAIN, S. R. and HOETING, J. A. (2010). A comparison study of extreme precipitation from six different regional climate models via spatial hierarchical modeling. *Extremes* **13** 219–239. [MR2643558](#)
- SMITH, R. L., TEBALDI, C., NYCHKA, D. and MEARNs, L. O. (2009). Bayesian modeling of uncertainty in ensembles of climate models. *J. Amer. Statist. Assoc.* **104** 97–116.
- SOLOMON, S., QIN, D., MANNING, M., CHEN, Z., MARQUIS, M., AVERYT, K. B., TIGNOR, M. and MILLER, H. L., EDS. (2007). *Climate Change 2007: The Physical Science Basis: Working Group I Contribution to the Fourth Assessment Report of the IPCC*. Cambridge Univ. Press, Cambridge, UK and New York, NY, USA.
- TEBALDI, C. and SANSÓ, B. (2009). Joint projections of temperature and precipitation change from multiple climate models: A hierarchical Bayes approach. *J. Roy. Statist. Soc. Ser. A* **172** 83–106.
- TEBALDI, C., SMITH, R. L., NYCHKA, D. and MEARNs, L. O. (2005). Quantifying uncertainty in projections of regional climate change: A Bayesian approach to the analysis of multimodel ensembles. *Journal of Climate* **18** 1524–1540.
- VER HOEF, J. M., CRESSIE, N. and BARRY, R. M. (2004). Flexible spatial models for kriging and cokriging using moving averages and the Fast Fourier Transform (FFT). *J. Comput. Graph. Statist.* **13** 265–282. [MR2063985](#)
- WHITTAKER, J. (1990). *Graphical Models in Applied Multivariate Statistics*. Wiley, New York. [MR1112133](#)

S. R. SAIN
 GEOPHYSICAL STATISTICS PROJECT
 INSTITUTE FOR MATHEMATICS APPLIED TO GEOSCIENCES
 NATIONAL CENTER FOR ATMOSPHERIC RESEARCH
 P.O. BOX 3000, BOULDER
 COLORADO 80307
 USA
 E-MAIL: ssain@ucar.edu

R. FURRER
 INSTITUTE OF MATHEMATICS
 UNIVERSITY OF ZURICH
 ZURICH
 SWITZERLAND
 E-MAIL: reinhard.furrer@math.uzh.ch

N. CRESSIE
 DEPARTMENT OF STATISTICS
 OHIO STATE UNIVERSITY
 COLUMBUS, OHIO 43210
 USA
 E-MAIL: ncressie@stat.osu.edu